

Performance analysis with atop

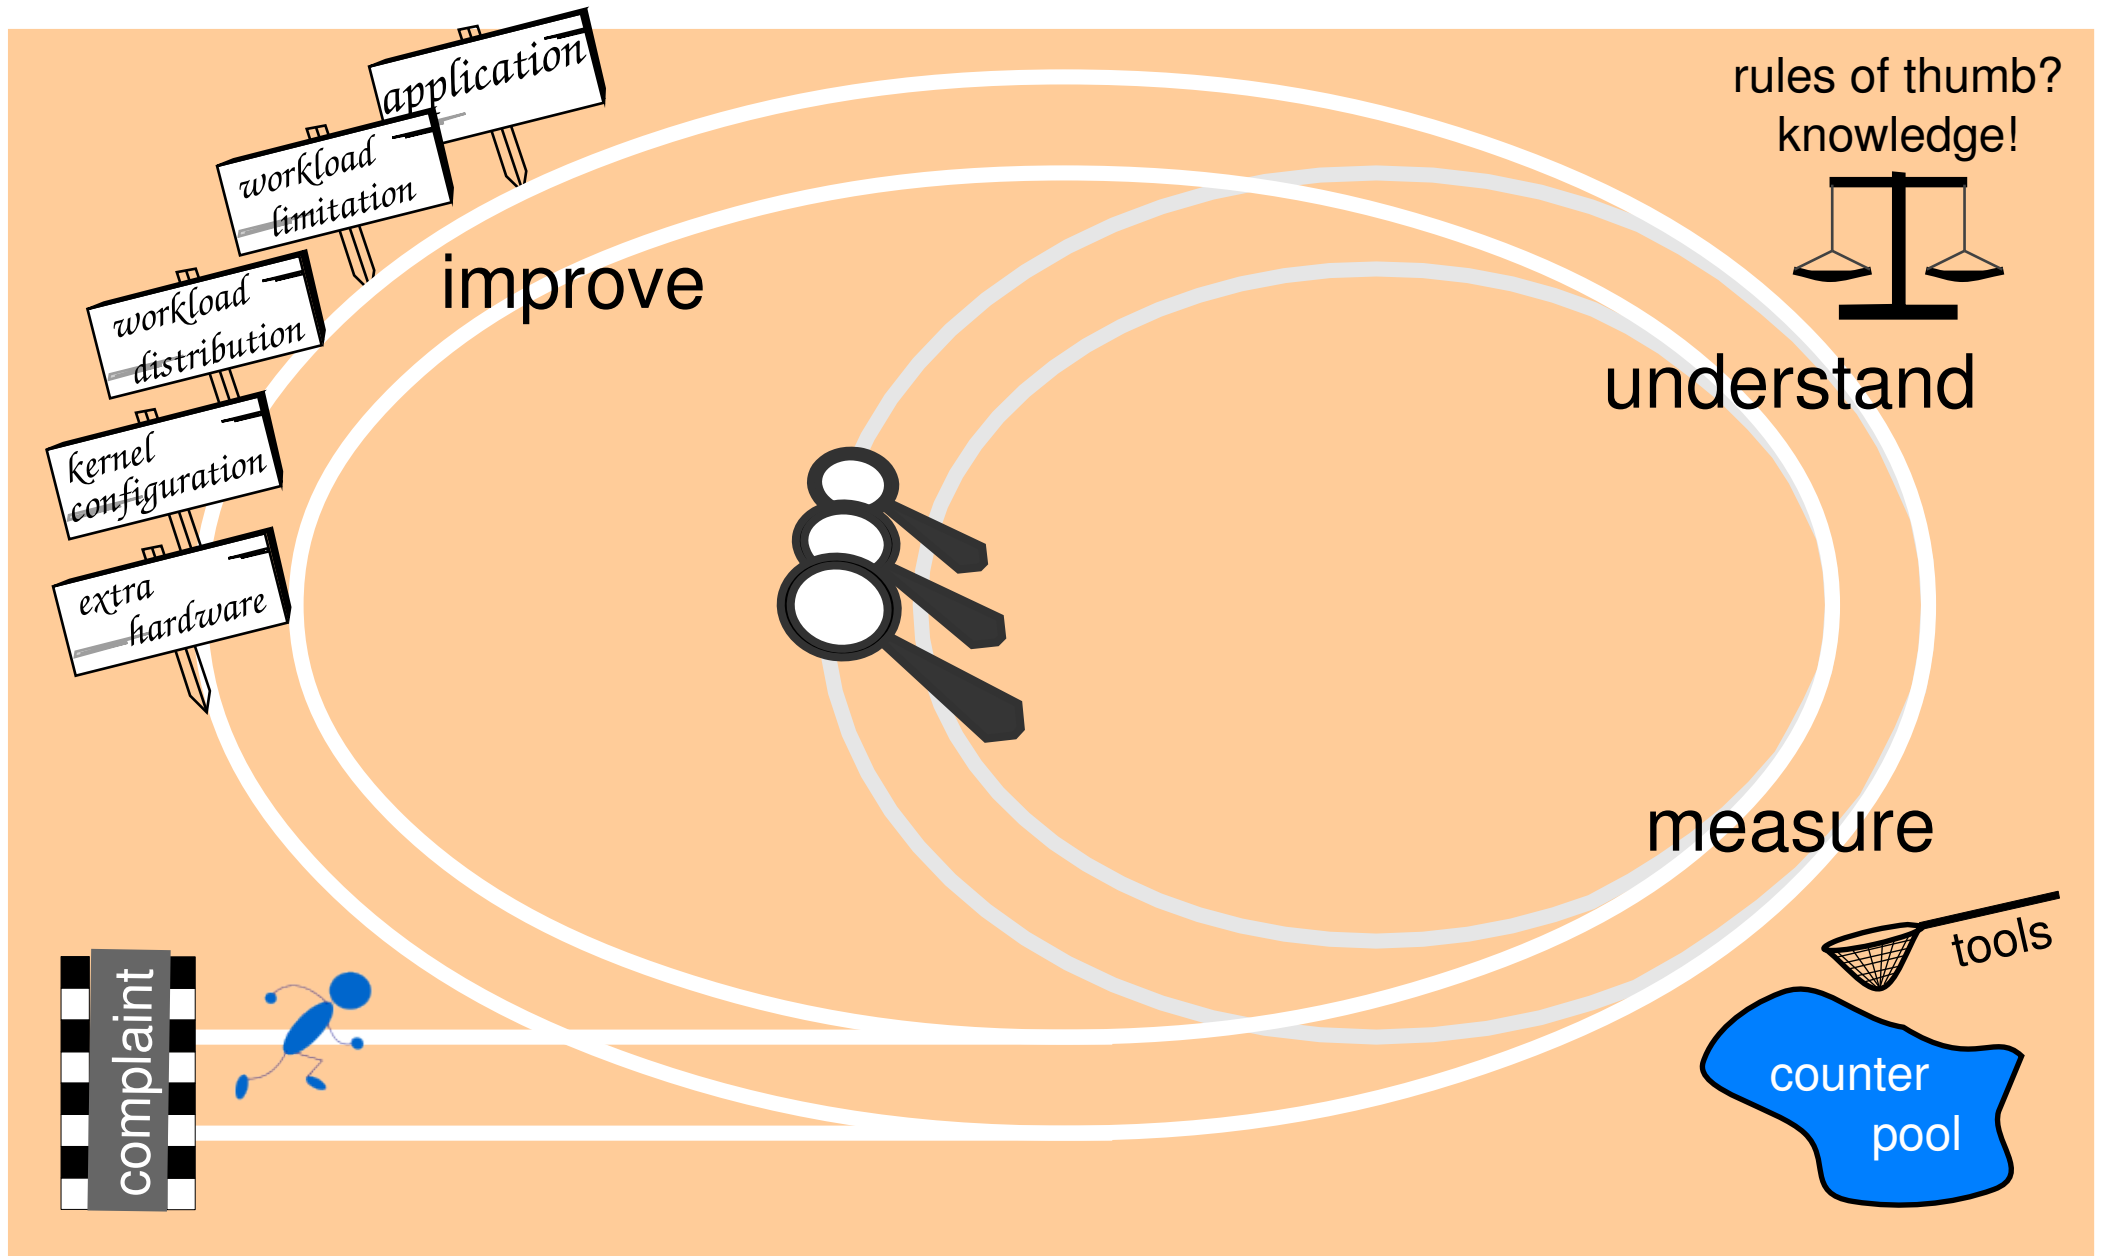


Gerlof Langeveld
gerlof@atoptool.nl



www.atcomputing.nl

Performance analysis — complaint till improvement



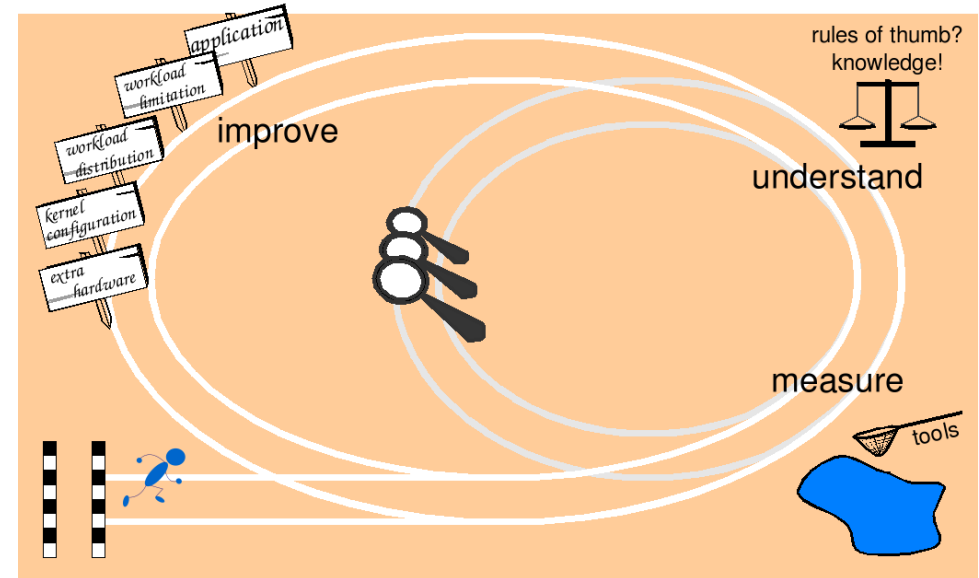
Complaint

Reason for analysis

- interactive use:
bad response times (seconds)
- batch processing:
throughput (transactions/second)

Nature of complaint

- structural
 - system *always* slow
 - analysis easy (at any moment)
- accidental
 - system *sometimes* slow
 - analysis of historical data (logging necessary)



Potential bottlenecks

Overload of certain hardware resources

- processor(s)

rule of thumb^{*)}: < 80% per CPU

- memory

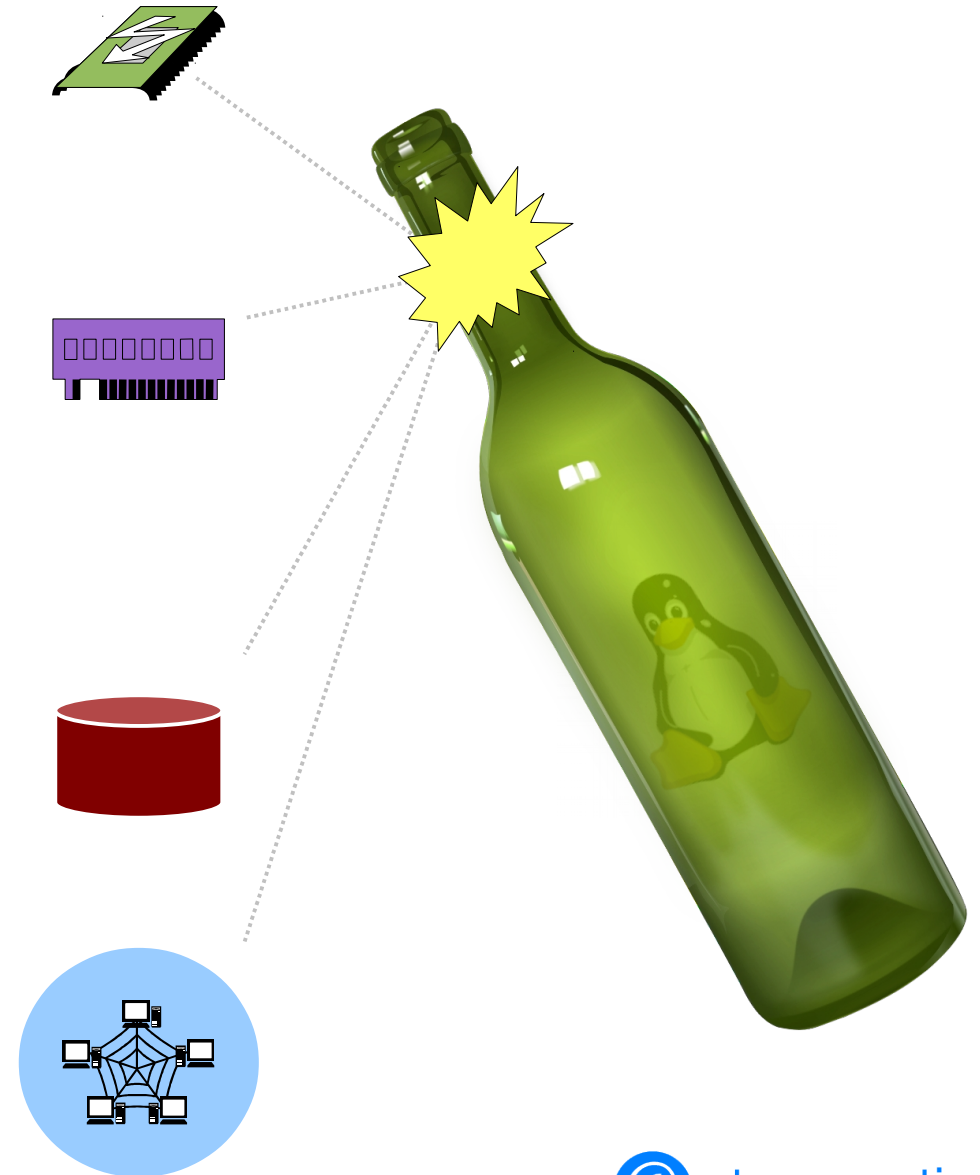
rule of thumb^{*)}: limited swapout
(preferably none)

- disk(s)

rule of thumb^{*)}: < 65% per disk

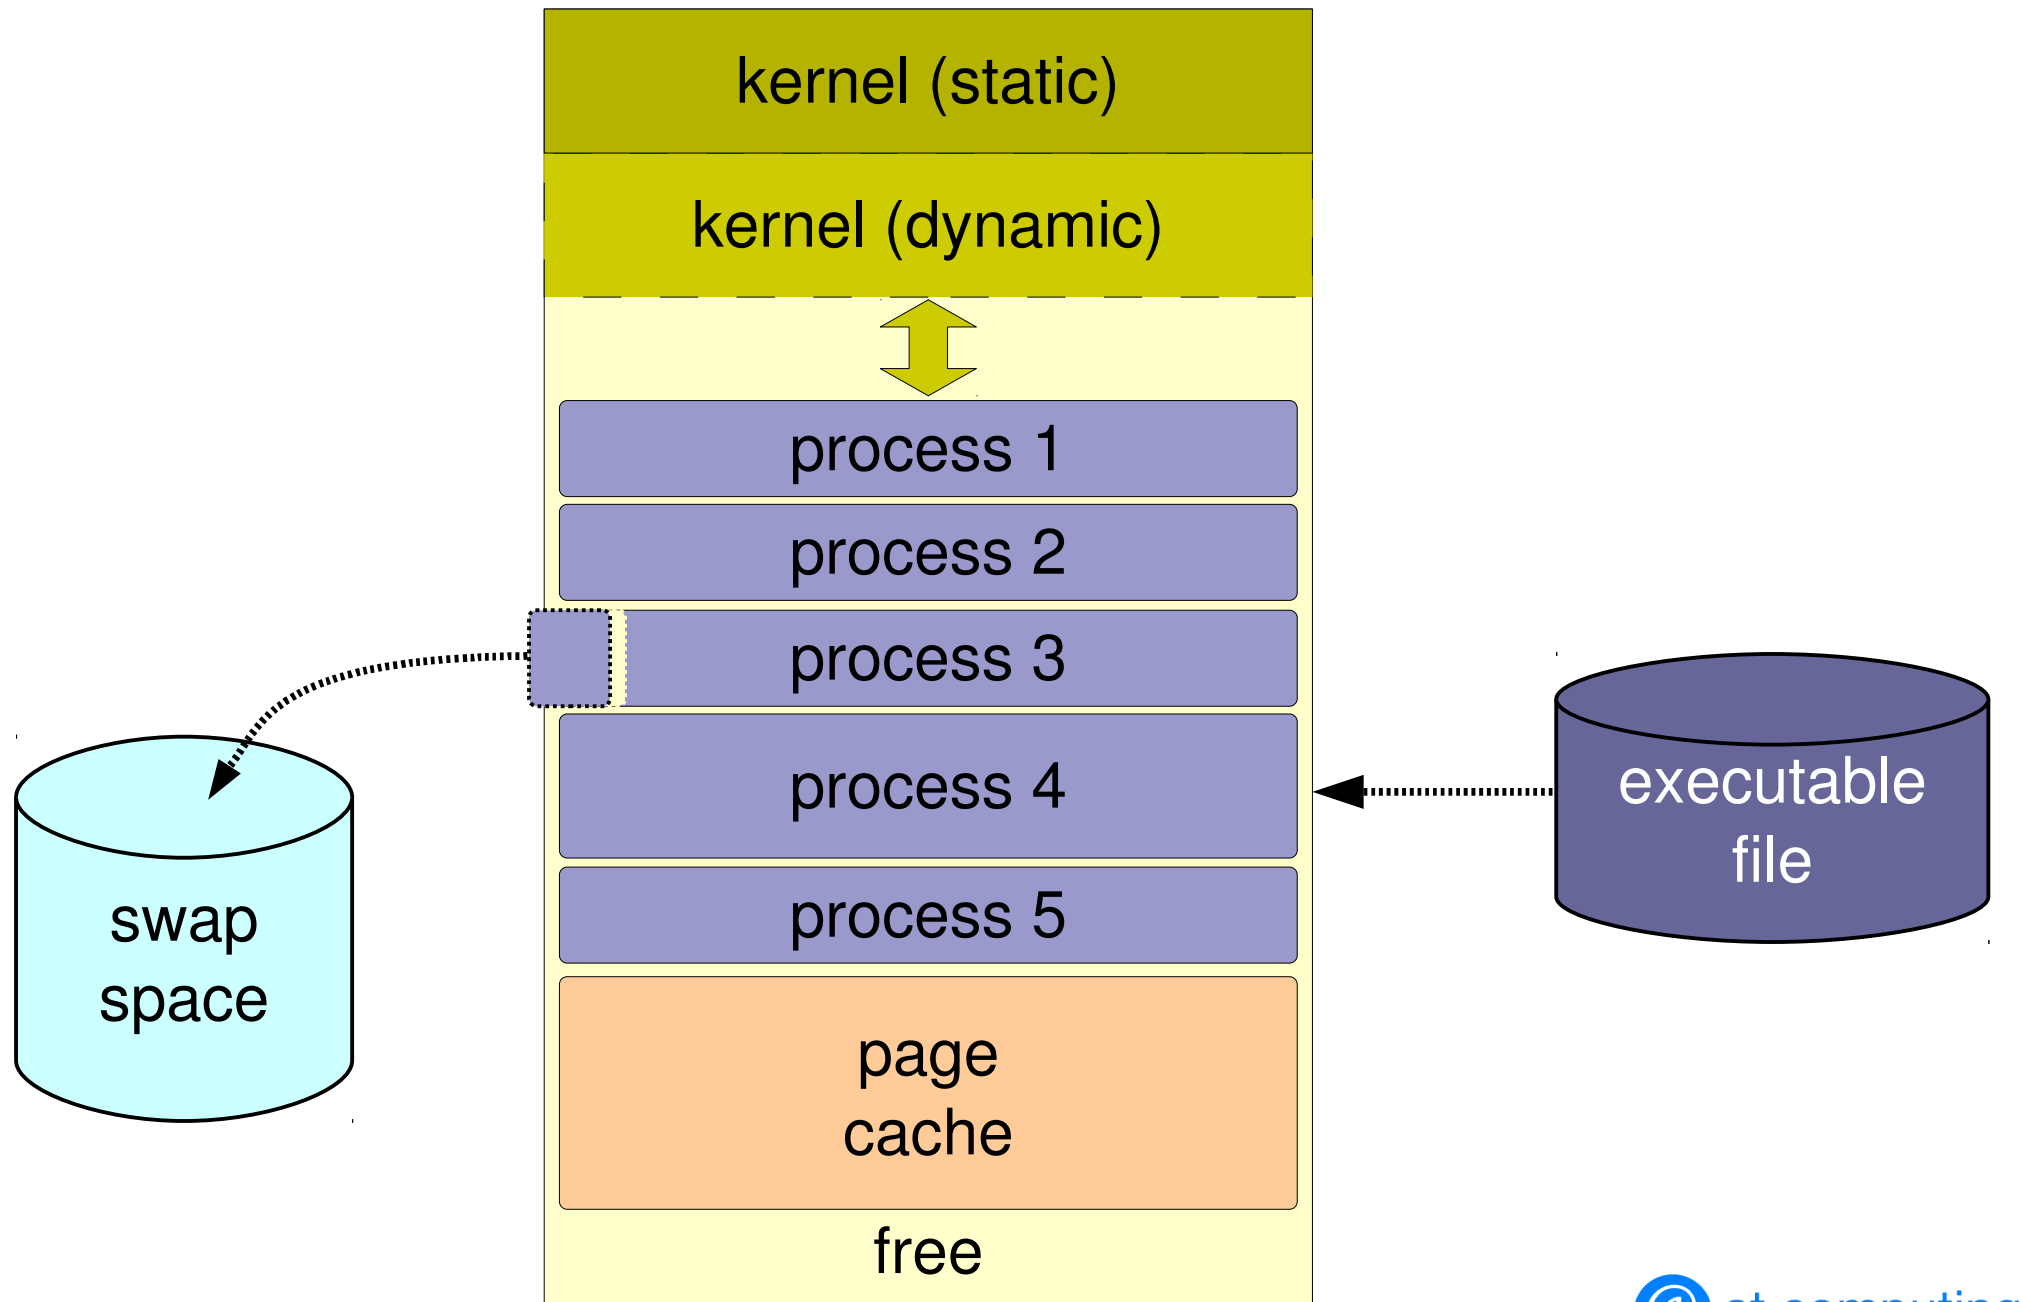
- network interface(s)

rule of thumb^{*)}: ?



^{*)} Beware of rules of thumb!

Memory management (simplified)



Overview generic measurement tools

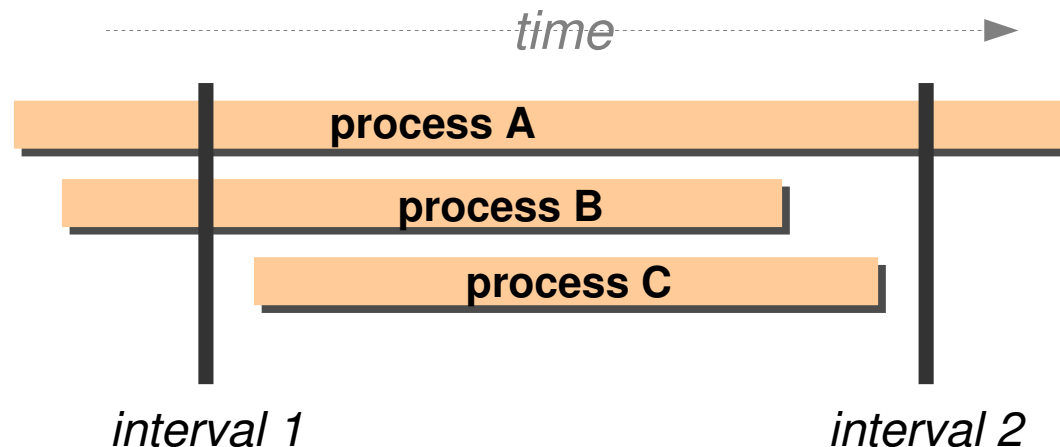
Generic tools	sar	vmstat	iostat	top	atop
Live measurements	✓	✓	✓	✓	✓
Historical data (logging)	✓	✗	✗	✗	✓
System-level data					
Processor	✓	✓	✓	✓	✓
Memory	✓	✓	✗	✓	✓
Disk	✓	✗	✓	✗	✓
Network	✓	✗	✗	✗	✓
Process-level data					
Processor	✗	✗	✗	✓	✓
Memory	✗	✗	✗	✓	✓
Disk	✗	✗	✗	✗	✓
Network	✗	✗	✗	✗	✓ with <i>netatop</i> kernel module

Measure with 'atop' – features

Characteristics

- utilization of *all* relevant hardware resources
 - CPUs
 - memory and swap space
 - disks and logical volumes (LVM)
 - network layers (including NFS)
- resource utilization of *all* processes even when processes have terminated (via process accounting)

} colors in case of overload



Measure with 'atop' – features

Characteristics

- data recording
 - permanent recording (default interval: 10 minutes)

<pre>\$ atop -r</pre>	today's data since midnight
<pre>\$ atop -r 20170223</pre>	data of specific date
<pre>\$ atop -r yy</pre>	data of day before yesterday

- branch to interval (b)
- next interval (t)
- previous interval (T)

- incidental recording

<pre>\$ atop -w /tmp/trial.atop 60 10</pre>	record 10 intervals of 60 seconds
<pre>\$ atop -r /tmp/trial.atop</pre>	data of specific date

Measure with 'atop' – features

Characteristics

- dynamic scaling of columns
- per interval only show
 - *utilized* resources (*all* resources: **f**)
 - *active* processes (*all* processes: **a**)
- show individual threads for multi-threaded processes (**y**)
- select process/thread information
 - generic – default (**g**) various (**v**) disk (**d**)
 - scheduling (**s**) memory (**m**) network (**n**)
- accumulate utilization of processes
 - per user (**u**) per container (**j**)
 - per program (**p**)


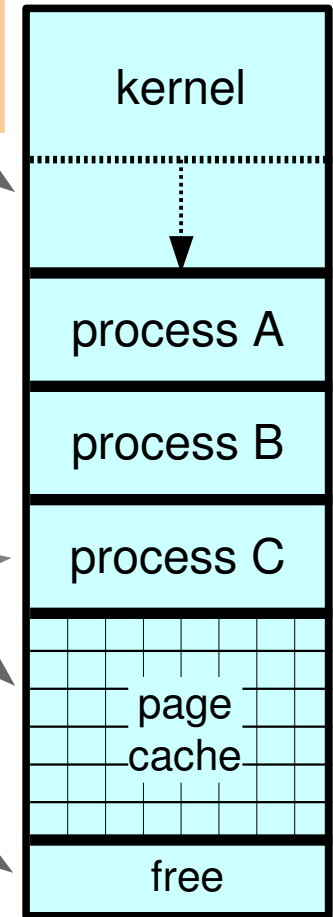
Measure with 'atop' – memory analysis

MEM	tot	993.8M	free	5.6M	cache	846M	buff	4.8M	slab	32.4M
SWP	tot	2.0G	free	2.0G			vmcom	398.3M	vmlim	2.5G
PAG	scan	22440	stall	55			swin	2	swout	3

PID	SYSCPU	USRCPU	VGROW	RGROW	USERNAME	THR	ST	EXC	S	CPU	CMD
30484	0.46s	7.02s	1548K	1548K	gerlof	1	--	-	R	75%	grep
27238	0.03s	0.27s	0K	4K	wwwadm	5	--	-	R	3%	firefox-
30555	0.01s	0.15s	8640K	5616K	root	1	N-	-	D	2%	mrtg
27187	0.08s	0.07s	0K	0K	wwwadm	1	--	-	R	2%	sshd
43	0.04s	0.00s	0K	0K	root	1	--	-	S	0%	kswapd0
27511	0.00s	0.03s	0K	0K	gerlof	4	--	-	R	0%	soffice.
27470	0.01s	0.02s	0K	0K	gerlof	1	--	-	S	0%	sshd
30198	0.01s	0.01s	0K	0K	gerlof	1	--	-	R	0%	atop
17963	0.01s	0.00s	0K	0K	apache	1	--	-	S	0%	httpd
30557	0.00s	0.01s	4088K	1844K	mailman	1	N-	-	D	0%	python
27251	0.00s	0.00s	0K	-12K	wwwadm	1	--	-	S	0%	gvim
17964	0.00s	0.00s	0K	0K	apache	1	--	-	S	0%	httpd
17919	0.00s	0.00s	0K	0K	root	1	--	-	S	0%	httpd
.....											

Keystroke m shows:

PID	MINFLT	MAJFLT	VSTEXT	VSIZE	RSIZE	VGROW	RGROW	MEM	CMD
27511	0	0	242K	246.3M	89788K	0K	0K	9%	soffice.
27238	1	1	9516K	98.3M	32884K	0K	4K	3%	firefox-
27251	0	0	2054K	19456K	10240K	0K	-12K	1%	gvim
17963	2	0	249K	21092K	6312K	0K	0K	1%	httpd
17964	3	0	249K	20964K	6236K	0K	0K	1%	httpd
.....									



Measure with 'atop' – disk analysis

LVM	vg00-lvusr	busy	95%	read	6668	write	2592	avio	1.17 ms
LVM	vg00-lvhome	busy	2%	read	0	write	28	avio	9.21 ms
LVM	vg00-lvroot	busy	1%	read	12	write	15	avio	5.00 ms
LVM	vg00-lvvar	busy	1%	read	0	write	10	avio	10.9 ms
DSK	sda	busy	95%	read	5439	write	172	avio	1.94 ms

PID	SYSCPU	USRCPU	VGROW	RGROW	RDDSK	WRDSK	ST	EXC	S	DSK	CMD
14318	0.61s	0.04s	0K	0K	77132K	0K	--	-	D	95%	grep
1915	0.01s	0.00s	0K	0K	0K	3584K	--	-	D	4%	jbd2/dm-5-8
1959	0.00s	0.00s	0K	0K	0K	48K	--	-	S	0%	flush-253:4
1920	0.00s	0.00s	0K	0K	0K	28K	--	-	S	0%	jbd2/dm-2-8
2283	0.00s	0.00s	0K	0K	0K	4K	--	-	S	0%	NetworkManager
2749	0.00s	0.00s	0K	0K	0K	4K	--	-	S	0%	ksmtuned
494	0.00s	0.00s	0K	0K	0K	4K	--	-	S	0%	jbd2/dm-0-8
1905	0.00s	0.00s	0K	0K	0K	4K	--	-	S	0%	jbd2/dm-4-8
3771	0.03s	0.07s	0K	0K	0K	0K	--	-	S	0%	plugin-contain

....

Keystroke d shows:

PID	TID	RDDSK	WRDSK	WCANCL	DSK	CMD
14318	-	77132K	0K	0K	95%	grep
1915	-	0K	3584K	0K	4%	jbd2/dm-5-8
1959	-	0K	48K	0K	0%	flush-253:4
1920	-	0K	28K	0K	0%	jbd2/dm-2-8
2283	-	0K	4K	0K	0%	NetworkManager

.....

Measure with 'atop' – network analysis

```
NET | transport      | tcpi 133548 | tcpo 18457 | udp_i      2 | udpo      0 |
NET | network         | ipi 133562 | ipo 18457 | ipfrw      0 | deliv 133559 |
NET | eth0            75% | pcki 130078 | pcko 629402 | si 8818 Kbps | so 752 Mbps |
```

```
  PID  SYSCPU  USRCPU  VGROW  RGROW  RDDSK  WRDSK  RNET  SNET  S  NET  CMD
14387  2.19s  0.12s   0K     0K     0K     0K   13e4  17e3  S 100% attract
14326  0.10s  0.15s   0K     4K     0K     0K   3426  1903  S   0% ssh
10316  0.00s  0.00s   0K     0K     0K     0K    20    10  S   0% ssh
 7135  0.10s  0.39s   0K    516K   0K     0K    12     8  S   0% thunderbird-bi
10310  0.00s  0.00s   0K     0K     0K     0K     1     2  S   0% ssh
 3655  0.45s  0.83s  232K   104K   0K     0K     0     0  S   0% gnome-terminal
.....
```

Keystroke **n** (with kernel module 'netatop') shows:

```
  PID  TCPRCV  TCPSND  UDPRCV  UDPSND  BANDWI  BANDWO  NET  CMD
14387  130084  16558   0        0  7284 Kbps  750 Mbps  100% attract
14326   3426   1903   0        0  1520 Kbps  109 Kbps   0% ssh
10316    20    10    0        0    1 Kbps    1 Kbps   0% ssh
 7135    12     8    0        0    0 Kbps    0 Kbps   0% thunderbird-bi
10310     1     2    0        0    0 Kbps    0 Kbps   0% ssh
 3655     0     0    0        0    0 Kbps    0 Kbps   0% gnome-terminal
.....
```

Measure with 'atopisar' – features

Characteristics

- system reports

```
live:  atopisar [-flags..]  interval [ samples ]
past:  atopisar [-flags..]  [-r file|date|y... ]
```

- examples: live CPU utilization

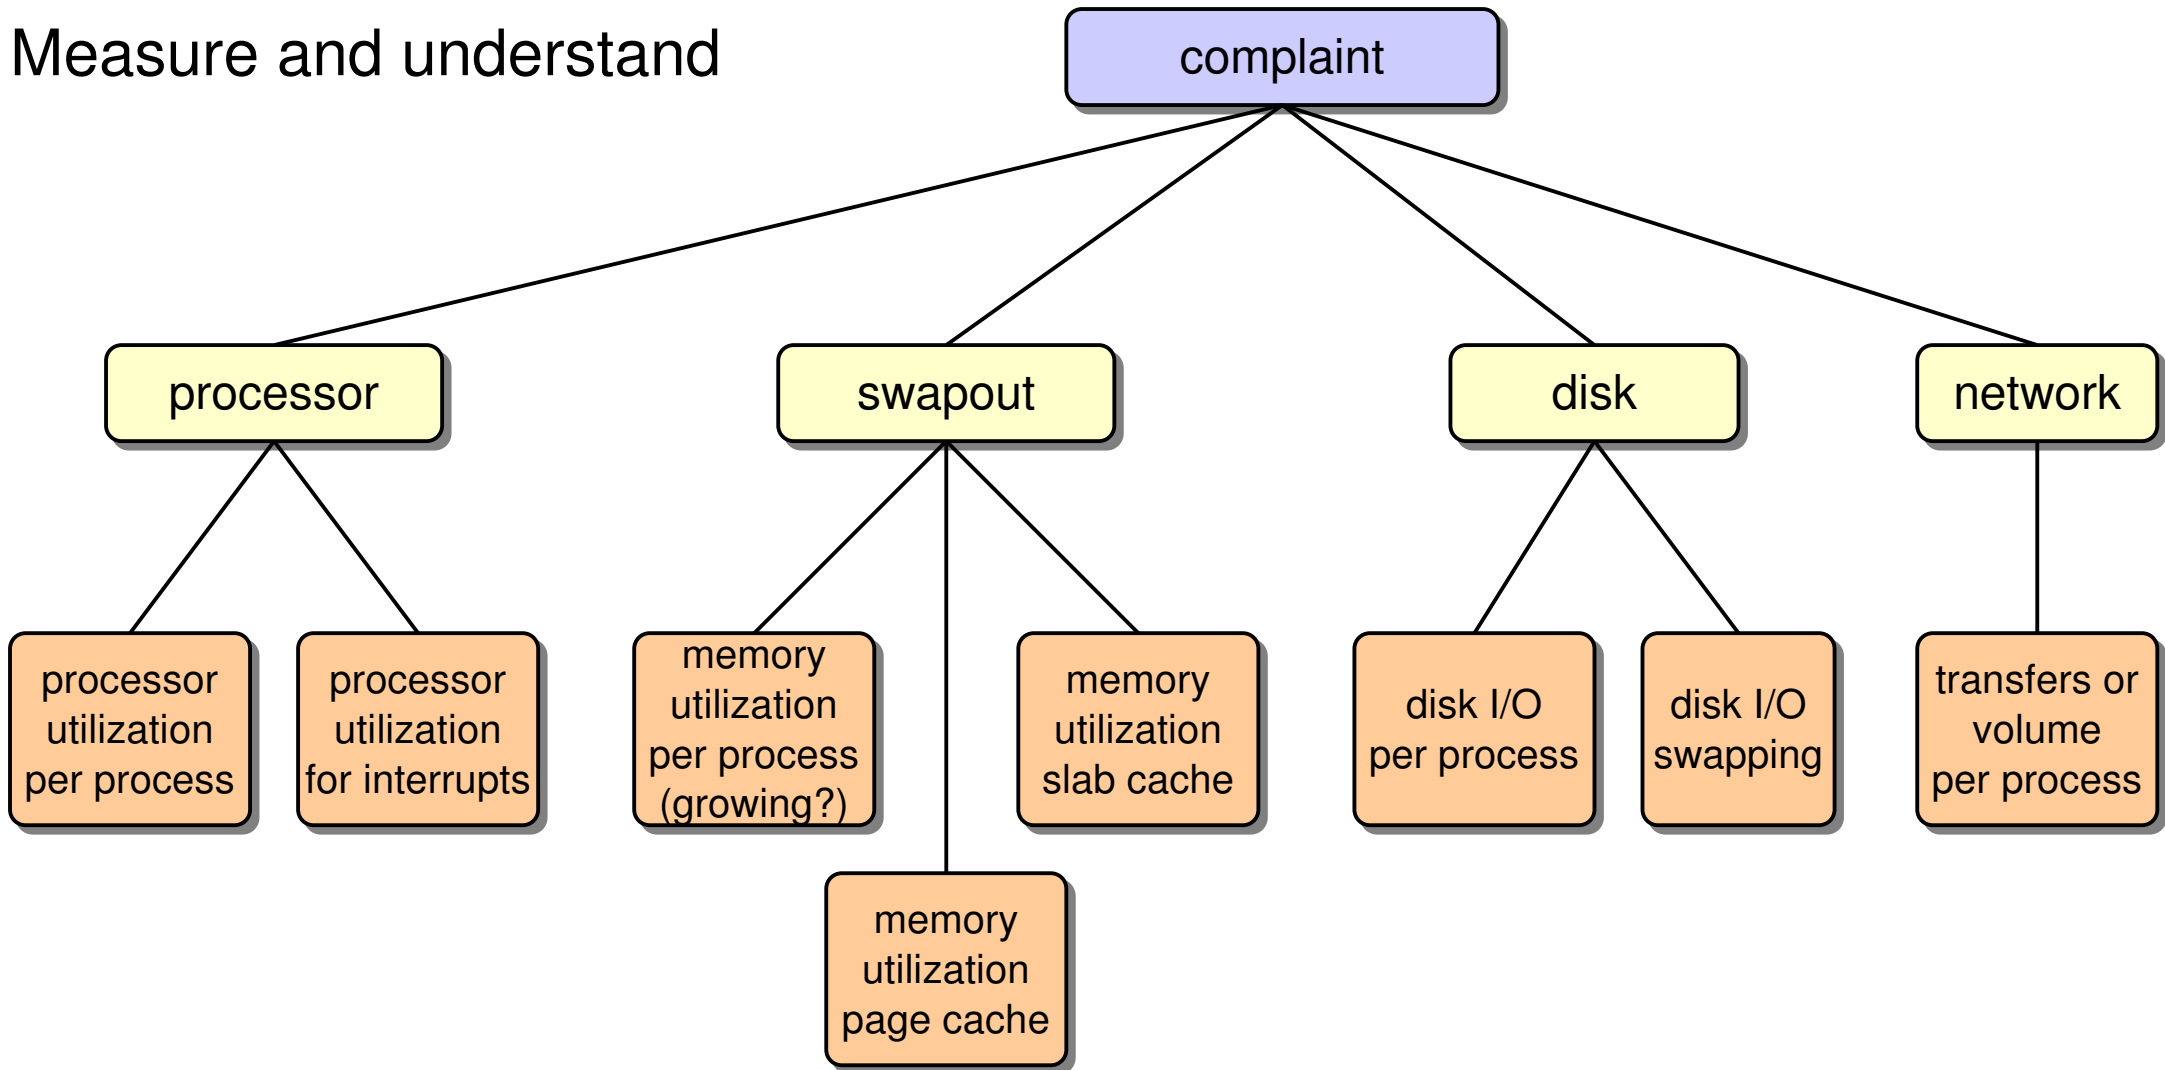
```
$ atopisar -c 60 5
myhost  3.10.0-118      ....      x86_64  2017/02/23

14:22:11  cpu %usr %nice  %sys %irq %softirq  %steal  %wait %idle
14:23:11  all  2    0    8    1    1    0    88    0
14:24:11  all  32   0   38   1    1    0    28    0
14:25:11  all  43   0   45   2    1    0    0    9
14:26:11  all  2    0    2    1    1    0    15   79
14:27:11  all  2    0    1    0    0    0    11   86
```

```
$ atopisar -A -r 20170223
$ atopisar -A -r yyyy
```

Measurement iteration

Measure and understand



Installation

Installation of **atop/atopsar**

- distribution repository
 - Ubuntu, Debian,
 - RedHat, Fedora, CentOS (from EPEL)
- website: www.atoptool.nl (RPM or tarball)
- GitHub: `git://github.com/atoptool/atop.git`



Installation of **netatop** kernel module

- website www.atoptool.nl (tarball)

Questions



?